ATLAS Computing Status

Simone Campana
Torre Wenaus
Davide Costanzo

ATLAS computing activity in Q2/Q3 2017

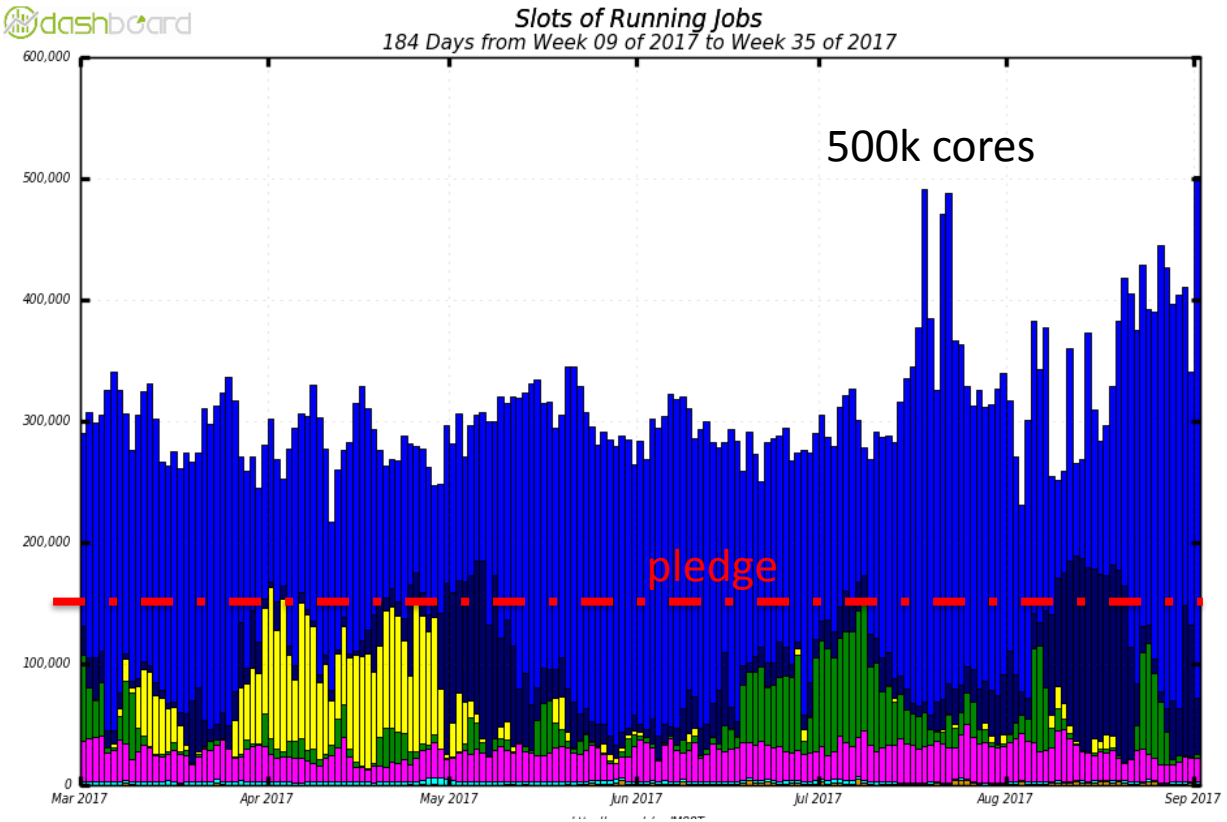

MC simulation and event generation

MC15 and MC16 campaign (more later)

MC reconstruction

User Analysis

Data reprocessing

2015+2016 data with R21 and final alignment

Analysis trains

On both data and MC

cores and core power

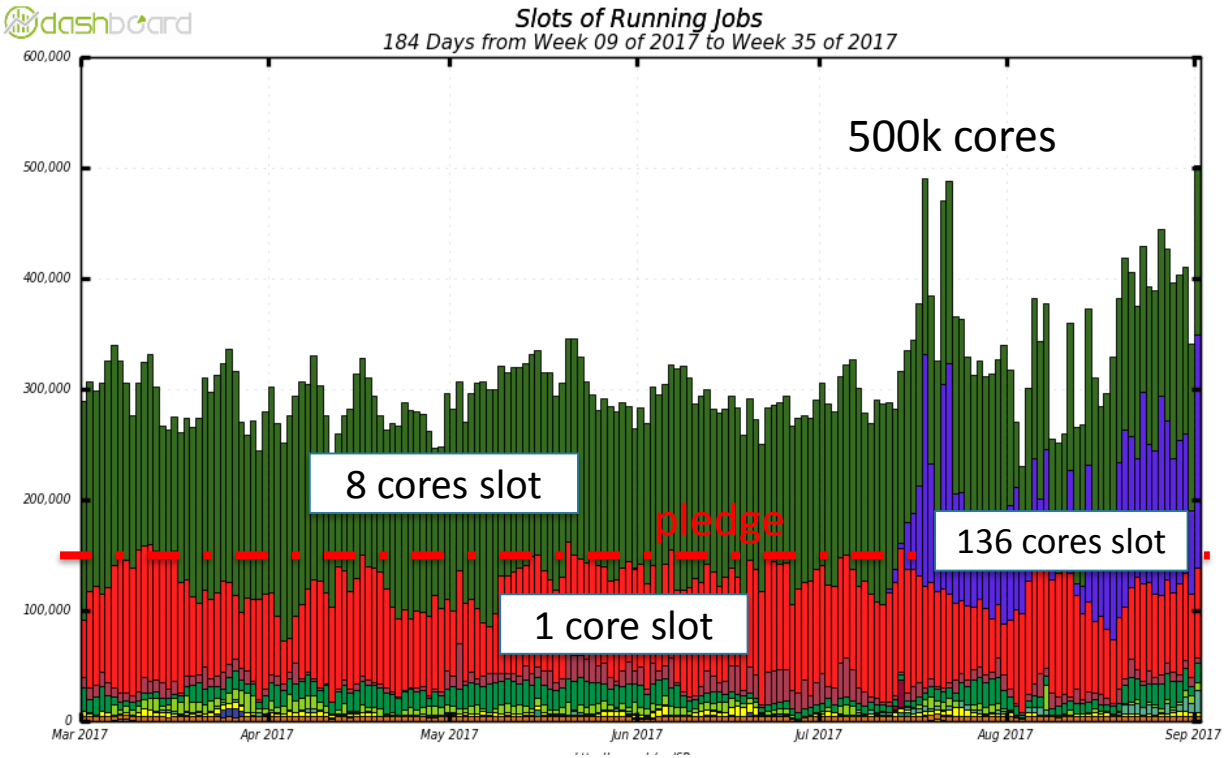

136 cores is CORI @ NERSC

1 core at CORI is 10x slower than an average Grid core

Looking at #core becomes more and more meaningless. The activity on benchmarking needs to be finalized

Still, a great demonstration of the ATLAS Workload Management System capabilities

Monte Carlo campaigns in 2017 and beyond

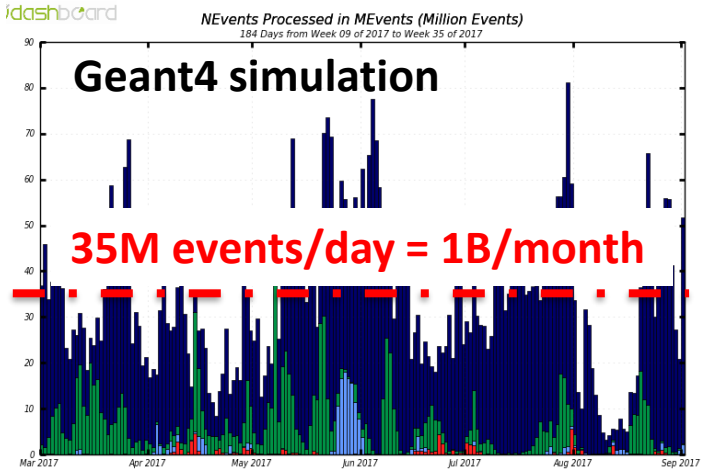

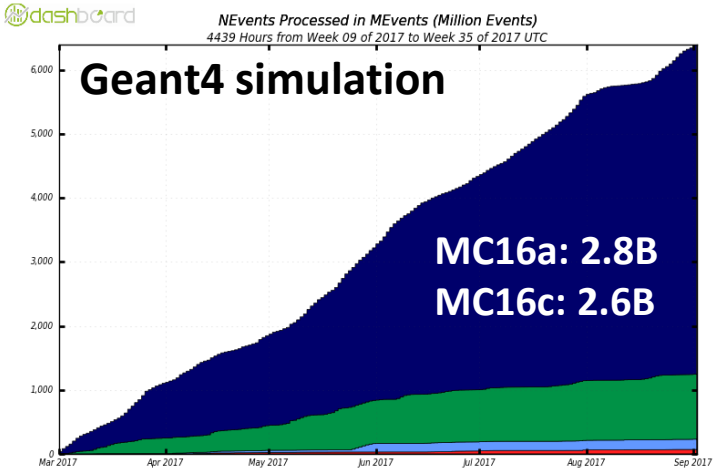

MC15: analyses based on 2015+2016 data only. Reconstructed with Athena 20.7

MC16: for 2015-2017 and the full Run-2 sample based analyses Reconstructed with Athena 21

Upgrade

MC16a matching 2015+2016 data with final conditions, trigger menu and pileup profile

(small) MC16b for trigger development at high pileup

MC16c matching 2017 data with estimated conditions

(?) MC16d: re-reconstruct MC16c HITS with final conditions and menu

MC16e matching 2018 data with estimated conditions

(?) MC16f: re-reconstruct MC16e HITS with final conditions and menu

Plans for 2017 and beyond

- We will reprocess the 2017 data with improved conditions
- (?) MC16d: re-reconstruct MC16c HITS with final conditions and menu. This will most likely happen, as the initially estimated MC pileup profile is so far considerably different with respect of the data
- MC15 is basically over. Continue with MC16

Campaign	Baseline	Alternative	Systematic	Sliced	Signal	Total
MC16a	3293 M	848 M	912 M	471 M	912 M	6436 M
MC16c	4116 M	1060 M	1140 M	589 M	1140 M	8046 M
MC16e	4116 M	1060 M	1140 M	589 M	1140 M	8046 M
Total	11525 M	2968 M	3192 M	1649 M	3192 M	**22528 M**

Table 1: The ongoing and forthcoming ATLAS Monte Carlo production campaigns for Run-2 analyses. Listed are the minimum required number of Monte Carlo events for each campaign and category.

Table from the document submitted to the CRSG: #events needed for "early full Run-2 analyses" (e.g. early searches) at 2019 summer conferences

1B evts/month: we will finish in March 2019

MC extensions in 2019 and 2020

Process	MC16a baseline #events	MC16 baseline #events	Extra Statistics factor	Extra Statistics #events
Inclusive Z	213 M	746 M	x3	2238 M
Inclusive W	200 M	700 M	x3	2100 M
V+jet	490 M	1715 M	x2	6860 M
Diboson	490 M	1716 M	None	0 M
Inclusive ttbar	200 M	700 M	x5	3500 M
ttbar+V	21	74	None	0 M
single top	20 M	70 M	x5	350 M
Multijet (x2 generators)	50 M	175 M	x5	1750 M
Total				**16798 M**

Table 2: The list of main MC16 samples organised by physics process. For most samples an increase in statistics will be needed in 2019 and 2020 to support precision measurements. The MC16 baseline statistic is calculated scaling the number of events for 2015+2016 data analyses by the increase of integrated luminosity. The extra statistics factor depends on the process and is defined based on experience with ongoing analyses and the observed impact of the lack of Monte Carlo statistics.

Table from the document submitted to the CRSG:
#events needed for the full exploitation of the Run-2 program, beyond "early full Run-2 analyses"

Extensions are calculated taking into account the expected 2017 and 2018 luminosity

The "Extra Statistics" factor is based on the current experience with precision measurements and the known statistical limitations

2019+2020 Monte Carlo needs: 16.8 B (extensions) + 6.0 B (new generators) + 3.5 B (Run-3 preparation) + 3.9 B (legacy from 2017/2018) = **30.2 B events** (8.5 B events Fast Sim and 21.7 Full Sim). We target to be able to produce at least 50% of these events in 2019.

CPU/WallTime efficiency

Follow up from the last RRB: loss of efficiency over the last several years. A complicated story.

Introducing MCORE caused a loss of 10% efficiency. Some of this is real, due to serial operations in MCORE environment (e.g. I/O). Some is artificial (initialization accounted differently in MCORE and SCORE)

EGI accounting

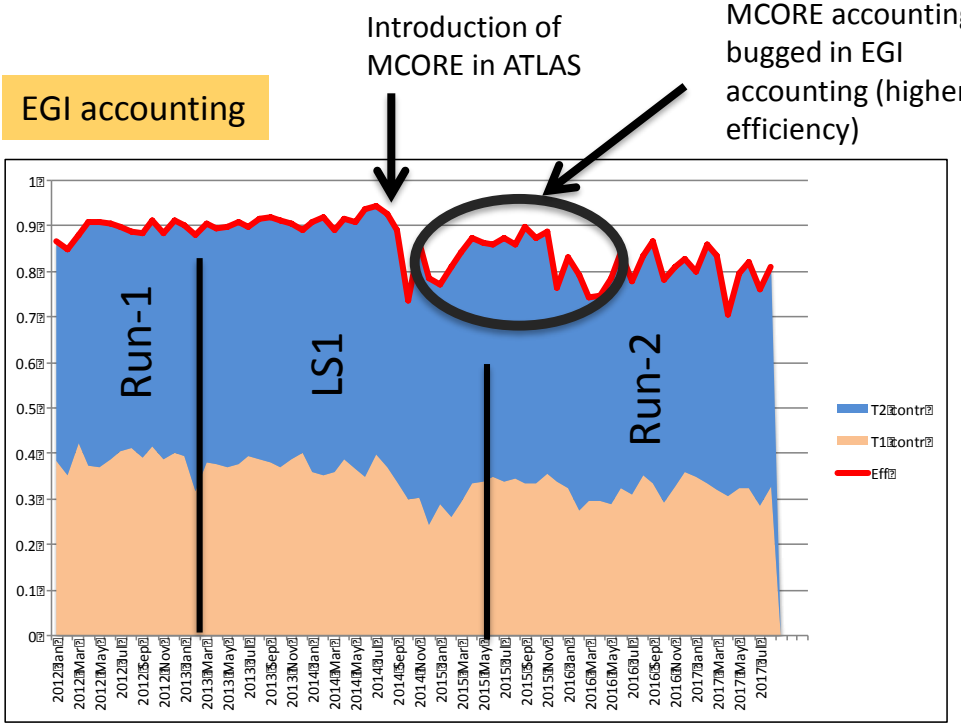

ATLAS dashboard – Run-2 period

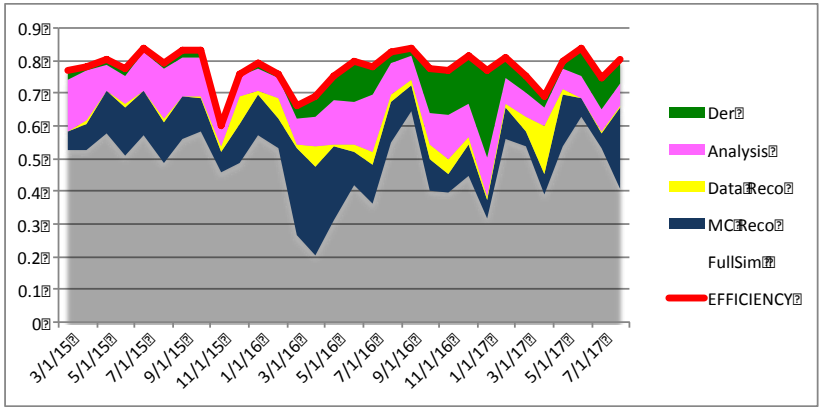

Conclusion-I: no obvious drop of ATLAS efficiency in Run 2

Conclusion-II: that does not mean we should accept 80% is good. We continue our work understanding performance and improving it

e.g. using checkpointed images to reduce serial init

Draining Loss when allocating MCORE slots

Following up from the discussion at the last RRB, we investigated with representative sites the “hidden” CPU loss in draining cores for MCORE allocation

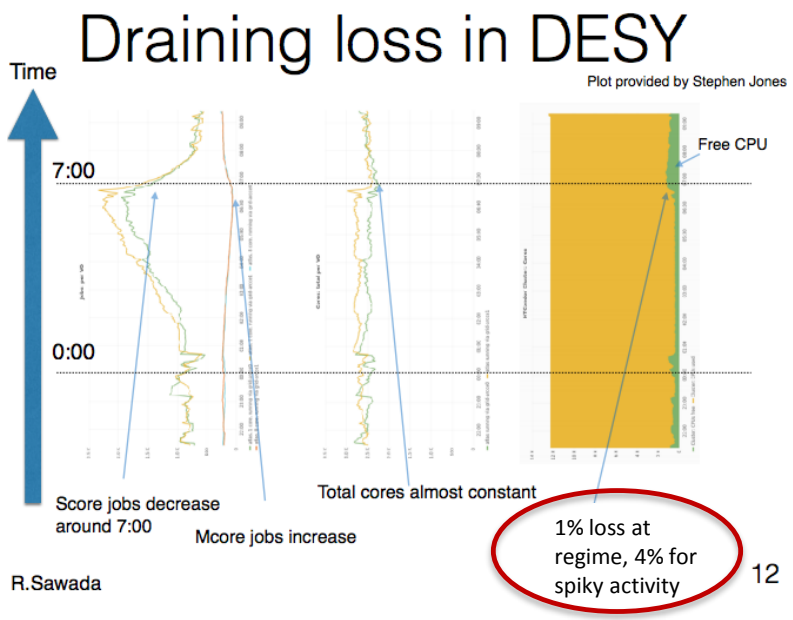

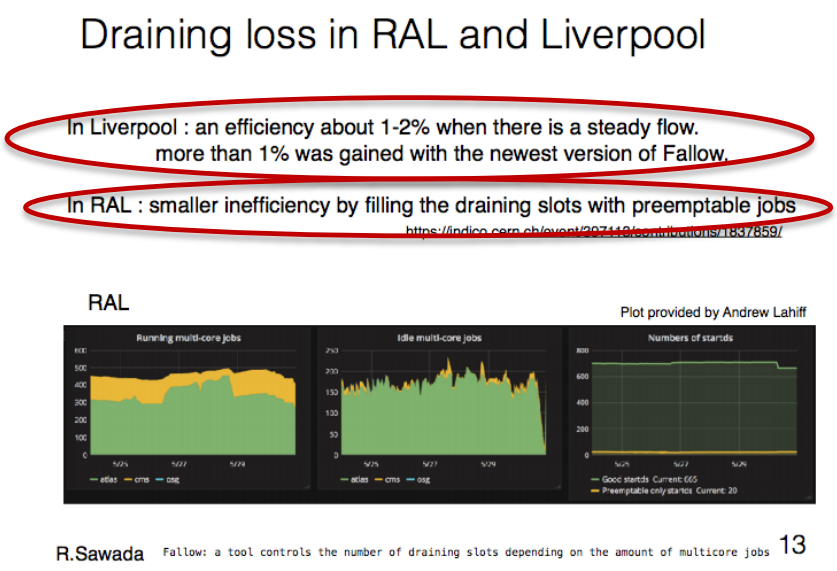

Conclusion: for ATLAS there is no hidden efficiency loss (< 1%). We fill 150% of our resources steadily with MCORE and steadily with SCORE, so no spikes.
When preemptible jobs are used, event service can ensure we make the most of their walltime.

Upgrading the Software toward Run-3

- The Athena Configuration system is being re-written
  - Among many benefits, alleviates some of the inefficiency mentioned before
- AthenaMT (Multi Threading) steadily progresses
  - We are largely done with the framework developments
  - Next step is to update the algorithmic code to drop thread-hostile constructs and use the thread-safe new framework features
  - Such a migration is a major effort, in many cases a rewrite. E.G for tracking the plan is to use ACTS
- Review of the ATLAS Conditions infrastructure in December
  - We review the proposed solution for Run-3
  - Participation of experts outside ATLAS (CMS, NA62, IT)

	Framework	Algorithms	
2015	Baseline Functionality	Basic demonstrators	
2016	Most functionality available ✓	Few more algos	
2017	All core functionality available	Start migration	
2018	Performance improvements	Bulk migration	← end of run 2
2019	bug fixes	Bulk migration	
2020	bug fixes	Validation	← start of run 3
2021	Run 3 Production		

Motivation ACTS start and philosophy

ACTS was branched off ATLAS Tracking about 2 years ago
- create a development area that allows more disruptive code changes:
  - *incorporating parallelism*
  - *review, modify and adapt Tracking code based on Run-1/2 experience*
- move to modern SW workflows (before ATLAS):
  git, CI, large coverage of Unit Tests
- create a longer-term perspective of code maintenance
  increase user and developer base (FCC, Tracking ML, …)

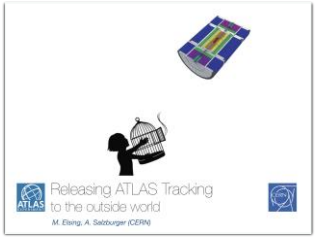

Re-integration into ATLAS/Athena was and remains the main goal of ACTS

Team ACTS group members/developers

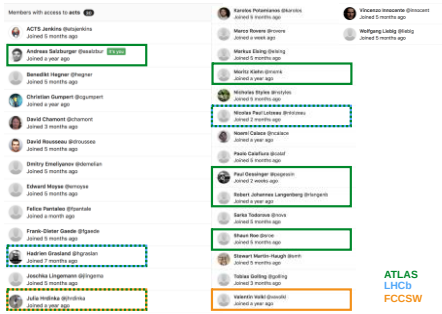

Input parameters, assumptions, disclaimers (ECFA/CHEP 2016)

Input Parameters at HL-LHC
(LOI = the ATLAS Letter of Intent for Upgrade Ph… -)

Output HLT rate: 10kHz (5 to 10 kHZ in LOI)
Reco time: 288s/event, Simul Time: 454 s/event at mu=200
Nr Events MC / Nr Events Data = 2
Fast Simulation: 50% of MC events
LHC live seconds /year: 5.5M

Simplified Computing M… …spec…
2016/2017 resource… …uests:

Data from previous yea… …t taken into a… …nt
=> Little difference at th… …inning of th… …n-4 but huge difference for Run-2 and H…

OLD

…ection of available …ources in HL-LHC:

20% more CPU/year
15% more storage/year

For the same cost

Projections evolve 2017 values
OF THIS SIMPLIFIED MODEL
(not the 2017 WLCG pledges)

Conclusion: looking at absolute numbers makes little sense.
Relative differences between needs and projections at HL-LHC are meaningful. With caveats.

HL-LHC baseline resource needs (ECFA/CHEP 2016)

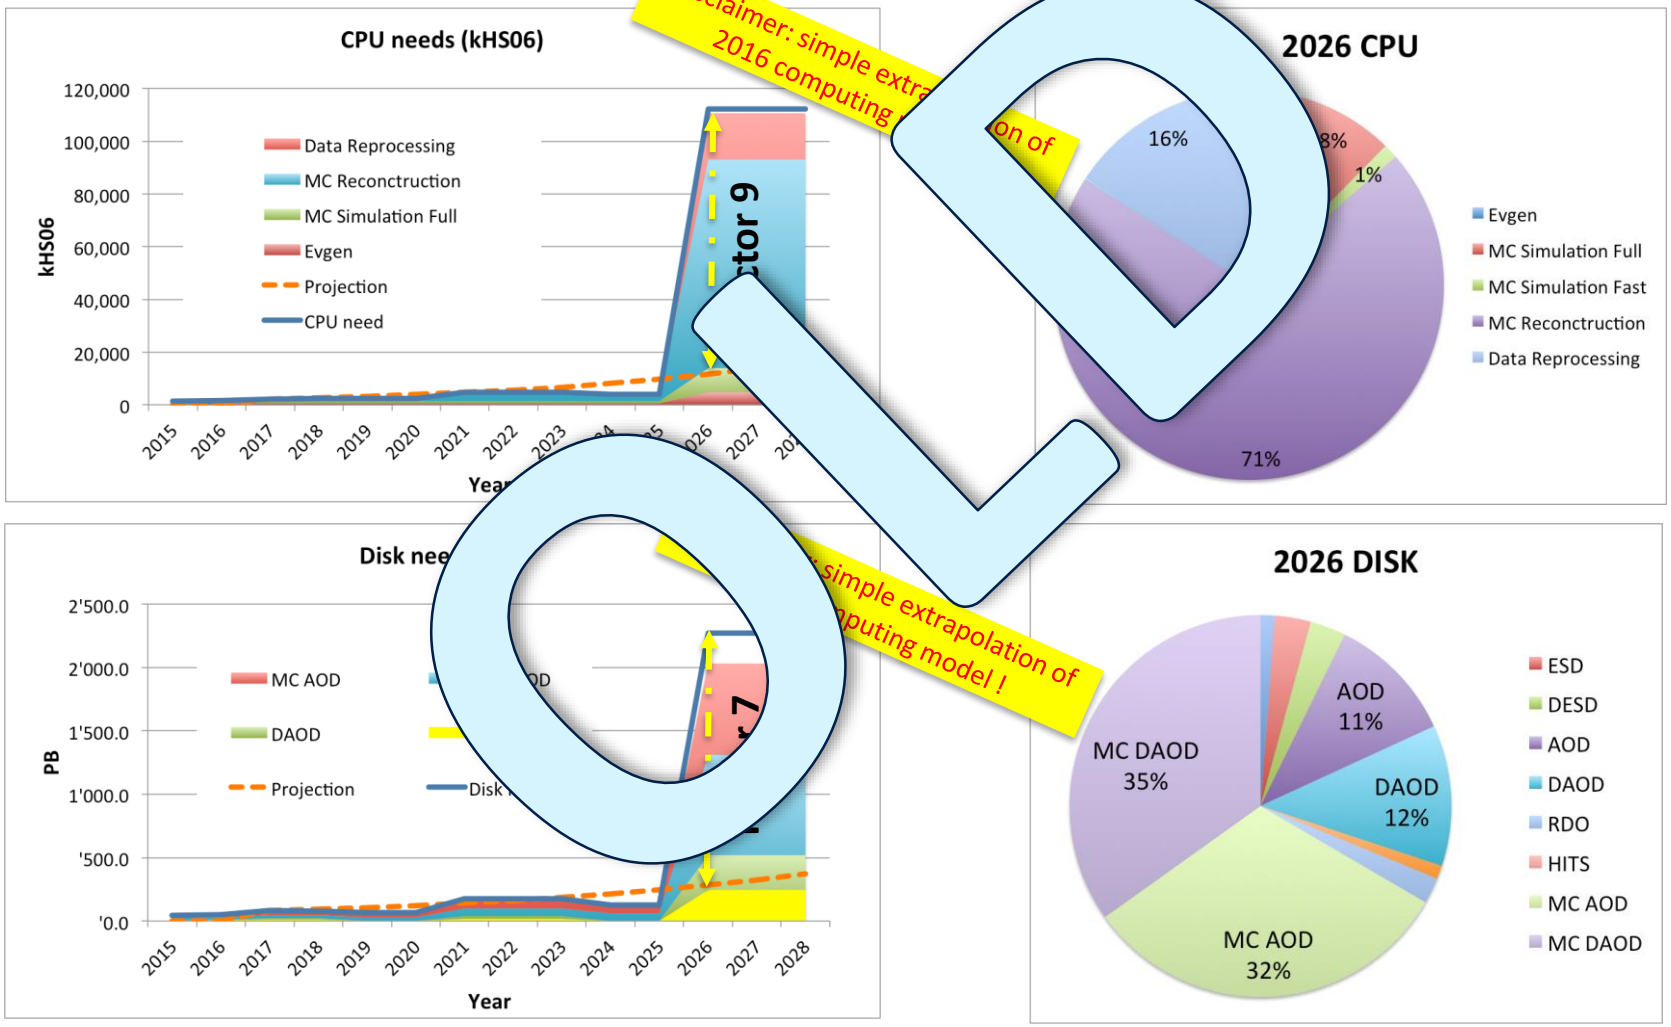

2017: Input parameters, assumptions, disclaimers

Input Parameters at HL-LHC,
updated after the conclusion of the Layout Task Force

Output HLT rate: 10kHz
Reco time: 130s/event at mu=200, Simul Time: 454 s/event
Nr Events MC / Nr Events Data = 1.5
N events with Fast Simulation: 50% of Full Simulation
LHC live seconds /year: 7.3 M

Less Simplified Computing Model with respect to ECFA 2016

Data from previous years **taken** into account
Tier-0 contribution added to the total

Projection of available resources in HL-LHC:

20% more CPU/year
15% more storage/year

For the same cost

Projections evolve 2017 values **OF THIS SIMPLIFIED MODEL (not the 2017 WLCG pledges)**

Conclusion: looking at absolute numbers makes little sense.
Relative differences between needs and projections at HL-LHC are meaningful. With caveats.

HL-LHC baseline resource needs (LHCC Sep. 2017)

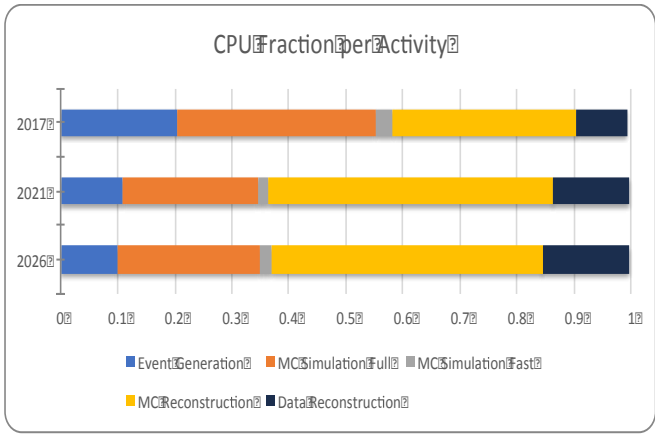

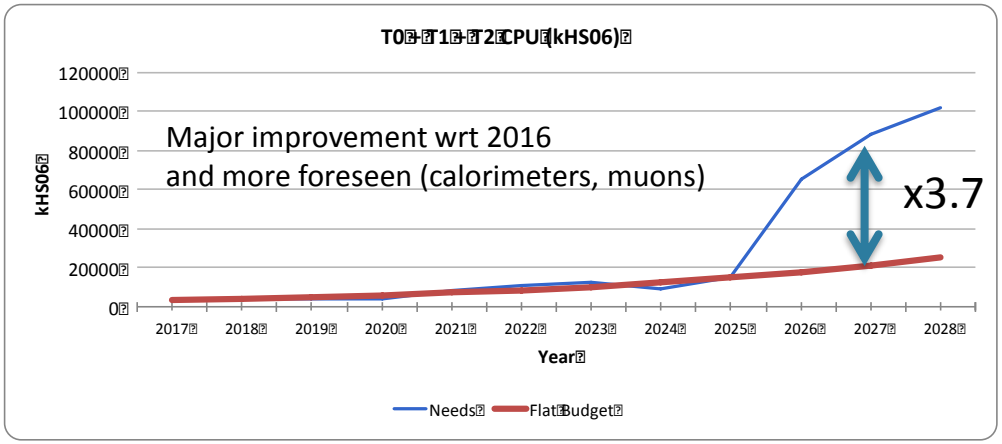

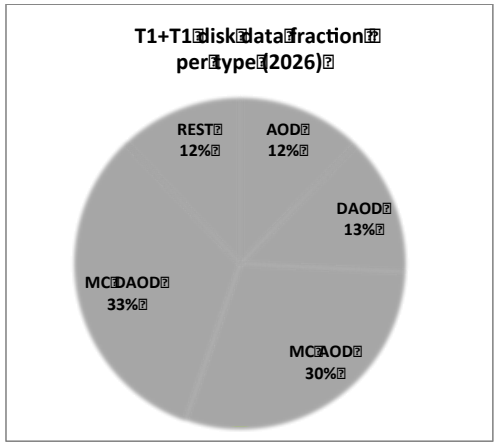

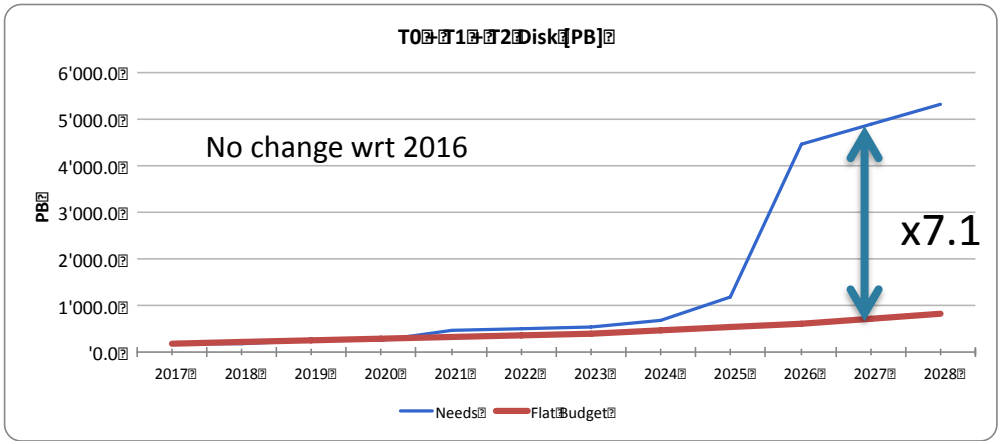

Conclusions and Announcement

Apologies for the longer-than-usual presentation but there were several topics to cover

It will not happen again. In fact this is my last meeting with the LHCC computing referees as ATLAS computing Coordinator. My term finishes on September 30

Torre Wenaus will take over as ATLAS Computing Coordinator with Davide Costanzo as deputy

I would like to thank the present and past referees, as well as the LHCC itself for the fruitful discussion and the support over the last 2.5 years